

AFRL-IF-RS-TR-2004-352
Final Technical Report
December 2004



COLLABORATIVE ASSISTANT FOR RAPID KNOWLEDGE FORMATION AND REASONING

George Mason University

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. K176

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2004-352 has been reviewed and is approved for publication

APPROVED: /s/

WILLIAM E. RZEPKA
Project Engineer

FOR THE DIRECTOR: /s/

JAMES A. COLLINS, Acting Chief
Advanced Computing Division
Information Directorate

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE DECEMBER 2004		3. REPORT TYPE AND DATES COVERED Final Apr 00 – Aug 03
4. TITLE AND SUBTITLE COLLABORATIVE ASSISTANT FOR RAPID KNOWLEDGE FORMATION AND REASONING			5. FUNDING NUMBERS C - F30602-00-2-0546 PE - 622301E PR - RKFM TA - 00 WU - 03	
6. AUTHOR(S) Gheorghe Tecuci				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) George Mason University MN 4A5 Learning Agents Center Computer Science Department 4400 University Drive Fairfax Virginia 22030-4444			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency AFRL/IFT 3701 North Fairfax Drive Arlington Virginia 22203-1714			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2004-352	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: William E. Rzepka/IFT/(315) 330-2762/ William.Rzepka@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) This research has developed a collaborative assistant for rapid knowledge formation and reasoning, called Disciple-RKF. Disciple-RKF enables a team of subject matter experts who do not have prior knowledge engineering experience to build an integrated knowledge base, with limited assistance from a knowledge engineer. Each subject matter expert teaches a personal Disciple-RKF agent while collaborating with it in solving specific problems. During this process, each Disciple-RKF agent learns from its expert and builds its knowledge base. All these knowledge bases are then merged into an integrated knowledge base for Disciple-RKF. Disciple-RKF has being successfully applied to develop knowledge-based agents for military center of gravity analysis which have been used in several courses at the US Army War College. Experimental results with Disciple-RKF demonstrate that the developed approach facilitates the acquisition of knowledge directly from subject matter experts and improves the knowledge base development process. The emphasis of this research was on acquiring expert's problem solving knowledge that is not normally represented in written documents. This task is complementary to that of acquiring knowledge that has already been expressed in textbooks or other documents.				
14. SUBJECT TERMS Artificial Intelligence, Knowledge Acquisition, Multistrategy Learning, Subject Matter Expert, Knowledge Engineering, Instructable Agent, Knowledge Base, Center Of Gravity Analysis				15. NUMBER OF PAGES 19
				16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

1. Introduction.....	1
2. Synergistic Collaboration and Transition to the Army War College.....	1
3. The Center of Gravity Analysis Challenge Problem	2
4. The Architecture of the Disciple-RKF Agents	3
5. Knowledge Base Development Methodology	5
6. Use of Disciple-COG in the “Case Studies of Center of Gravity Analysis” Courses: Scenario Elicitation Experiments.....	8
7. Use of Disciple-RKF/COG in the “Military Applications of Artificial Intelligence” Courses: Agent Teaching Experiments.....	8
8. Experiment of Parallel Knowledge Base Development by Subject Matter Experts.....	9
9. Conclusions and Future Research Directions	11
10. Appendix.....	12
10.1 Publications describing the results of this project.....	12

List of Figures

Figure 1: General Approach to Rapid Knowledge Formation by Subject Matter Experts.....	1
Figure 2: Synergistic integration of artificial intelligence research military strategy research, and military education.	2
Figure 3: The identification and testing approach to center of gravity analysis.	3
Figure 4: The Architecture of the Disciple-RKF Agents.....	5
Figure 5: Three ontology browsers of Disciple-RKF.	6
Figure 6: The reasoning tree and the modeling assistant interfaces.	7
Figure 7: Task formation and example explanation.	8
Figure 8: Learned Rule	9
Figure 9: Training and using the Disciple agent (the COG course).	9
Figure 10: Teaching and testing a personal Disciple agent (the MAAI courage).	10
Figure 11: Experiment of rapid knowledge base development by subject matter experts.	11

1. Introduction

The general objective of this research project was to develop and to experimentally validate a collaborative assistant for rapid knowledge formation and reasoning. This assistant will enable a team of Subject Matter Experts (SME) who do not have prior knowledge engineering experience to rapidly construct, update and extend an integrated knowledge base for a complex application. The emphasis of this research was on acquiring expert's problem solving knowledge that is not normally represented in written documents. This task is complementary to that of acquiring knowledge that has already been expressed in textbooks or other documents. This research has benefited from the basic research done under the AFOSR's Software and Systems project titled "Mixed-Initiative Knowledge Base Development" (2000-2003) and from collaboration with the US Army War College. It resulted in the development, experimental use, and transition of a complex knowledge engineering environment, called Disciple-Rapid Knowledge Formation (RKF), and its application to the military Center of Gravity (COG) analysis domain, as described in the following, and in the papers from section 10.1.




Figure 1: General Approach to Rapid Knowledge Formation by Subject Matter Experts

Figure 1 introduces the general approach investigated. In this approach, each subject matter expert teaches a personal Disciple-RKF agent, while collaborating with it in solving specific problems. During this process, the Disciple-RKF agent learns from the expert, building, extending and improving its knowledge base. The resulting knowledge bases of all these Disciple-RKF agents are then integrated by a knowledge engineer. The Disciple-RKF agent with the integrated knowledge base can then be used in three ways. It can be used by a non-expert as a problem solver. It can be used by an expert as a problem solving assistant. Finally, it can be used by a student as a tutoring system.

2. Synergistic Collaboration and Transition to the Army War College

This project involved a multi-objective collaboration between both the Learning Agents Center of George Mason University and the Center for Strategic Leadership along with the Department of Military Strategy, Planning, and Operations of the US Army War College (Tecuci et al., 2002a). The US Army War College provided the challenge problem (strategic center of gravity analysis), extensive subject matter expertise (faculty and high-ranking officers from all the military services), and experimentation support for both the developed technology and the resulting knowledge bases and agents.

A distinguishing feature of this collaboration is the synergistic integration of artificial intelligence research, with military strategy research, and the practical use of agents in education, as indicated in Figure 2.

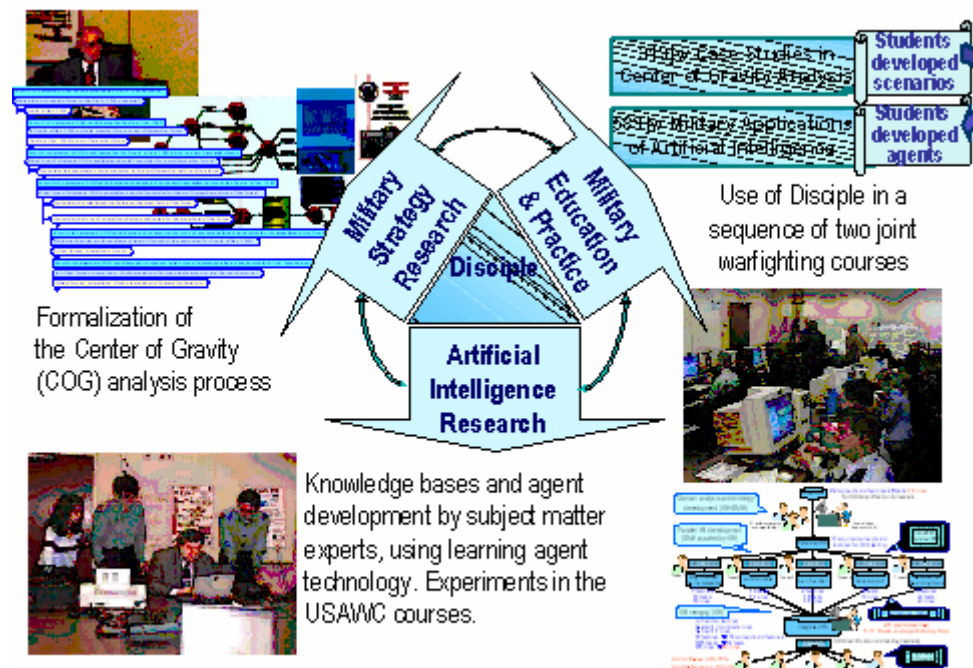


Figure 2: Synergistic integration of artificial intelligence research military strategy research, and military education.

The artificial intelligence research objective was the development of knowledge bases and agents by subject matter experts using learning agent technology. The military strategy research objective was the development of a systematic approach to center of gravity determination. The educational objective was the enhancement of the educational process of senior military officers and strategic leaders through the use of intelligent agent technology. This integration accelerated the development of the artificial intelligence technology and of intelligent agents for center of gravity analysis. It also facilitated the transition of this work to the US Army War College, where Disciple agents have been used since 2001 in a sequence of two joint warfighting courses, “319jw Case Studies in Center of Gravity Analysis,” and “589jw Military Applications of Artificial Intelligence.”

3. The Center of Gravity Analysis Challenge Problem

Military center of gravity analysis was used as a challenge problem to test the knowledge acquisition, learning, and problem solving methods of Disciple-RKF. The concept of center of gravity, introduced by Karl von Clausewitz (Clausewitz, K.V. 1832. *On War*, translated and edited by M. Howard and P. Paret. Princeton, NJ: Princeton University Press, 1976), is fundamental to military strategy, denoting the primary source of moral or physical strength, power, or resistance of a force (Strange, J. 1996. *Centers of Gravity & Critical Vulnerabilities: Building on the Clausewitzian Foundation So That We Can All Speak the Same Language*.

Quantico, Virginia, USA, Marine Corps University). The most important objective of a force (state, alliance, coalition, or group), in any type of conflict, is to protect its own center of gravity while attacking the center of gravity of its enemy. Therefore, in the education of strategic leaders at all U.S. senior military service colleges, there is great emphasis on center of gravity analysis. This analysis requires a wide range of background knowledge not only from the military domain, but also from the political, psychosocial, economic, geographic, demographic, historic, international, and other domains. In addition, the situation, the adversaries involved, their goals, and their capabilities can vary in important ways from one scenario to another. Center of gravity analysis is a very good example of knowledge-intensive, expert problem-solving that a Disciple agent should be able to learn.

Our approach to center of gravity analysis, based on the work of Strange (1996) and Giles and Galvin (Giles, P.K., and Galvin, T.P. 1996. *Center of Gravity: Determination, Analysis and Application*. CSL, U.S. Army War College, PA: Carlisle Barracks), and developed with experts from the US Army War College, consists of two main phases, *identification* and *testing*, as shown in Figure 3. During the identification phase, center of gravity candidates from different elements of power of a force (such as government, military, people, economy) are identified. For instance, a strong leader is a center of gravity candidate with respect to the government of a force. Then, during the testing phase, each candidate is analyzed to determine whether it has all the critical capabilities that are necessary to be the center of gravity. For example, a leader needs

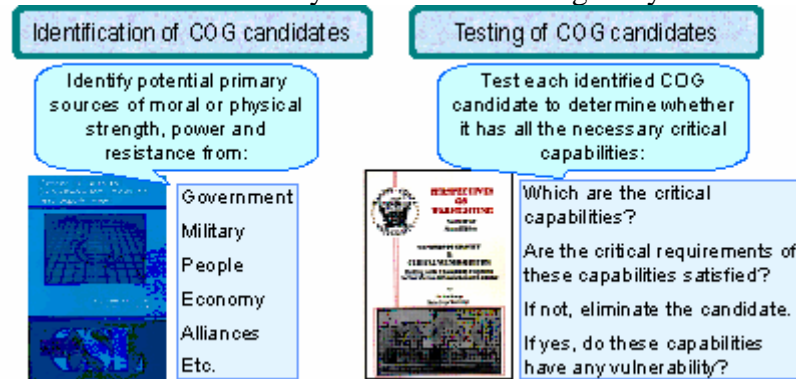


Figure 3: The identification and testing approach to center of gravity analysis.

to be protected, stay informed, communicate (with the government, the military, and the people), be influential (with the government, the military, and the people), be a driving force, have support (from the government, the military, and the people), and be irreplaceable. For each capability, one needs to determine the existence of the essential conditions, resources and means that are required by that capability to be fully operative, and which of these, if any, represent critical vulnerabilities.

4. The Architecture of the Disciple-RKF Agents

The architecture of Disciple-RKF includes the components from Figure 4 (Boicu et al., 2004). The core of the system is the learning agent shell, which has the following domain-independent components:

- A problem solving component based on the task reduction paradigm of problem solving, including the following modules:
 - Modeling assistant that helps the user to express his/her contributions to the problem solving process;
 - Interactive problem solving agent;
 - Autonomous problem solving agent.
- A knowledge acquisition and learning component for acquiring and refining the knowledge of the agent and allowing a wide range of operations (including ontology import, user definition of knowledge base elements, ontology learning, and rule learning). This component includes the following modules:
 - Ontology development modules:
 - Tree-based browsers for objects and features;
 - Graph-based browsers for objects and features (association browser, hierarchical browser);
 - Viewers and Editors for objects and features;
 - Ontology import module;
 - Knowledge base merging module.
 - Instances elicitation modules:
 - Scenario elicitation module;
 - Scripts editor.
 - Learning and refining modules:
 - Task formalization and learning module;
 - Explanation generation module;
 - Rule learning module;
 - Rule refinement with positive examples;
 - Rule refinement with negative examples;
 - Rule analysis module;
 - Rule regeneration module;
 - Exceptions-based knowledge base refinement;
 - Feature learning module.
- A knowledge base manager which controls the access to and the updates of the knowledge base. Each module of Disciple-RKF can access the knowledge base only through the functions of the knowledge base manager.
- A windows-based, domain-independent, graphical user interface.

The three components in the right hand side of Figure 4 are the typical domain dependent components of a Disciple-RKF agent that was customized for a specific application, such as center of gravity analysis. The Disciple-RKF/COG agent (also referred as Disciple-COG) includes the following components:

- A customized problem solving component that extends the basic task-reduction component in order to satisfy the specific problem solving requirements of the application domain.

- Customized graphical user interfaces which are built for the specific Disciple agent to allow the experts and the end users to communicate with the agent as close as possible to the way they communicate in their domains.
- The knowledge base of the Disciple agent that contains knowledge specific to the center of gravity analysis domain.

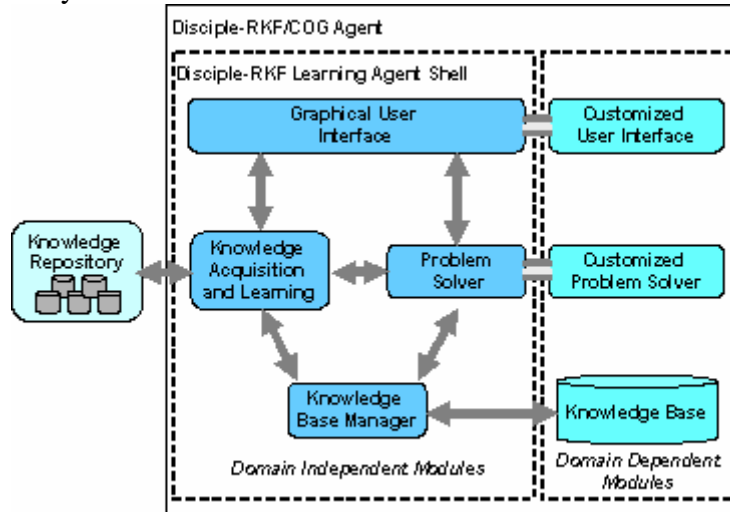


Figure 4: The Architecture of the Disciple-RKF Agents

5. Knowledge Base Development Methodology

The Disciple approach covers all the phases of agent development and use (Boicu et al., 2001c; Tecuci et al., 2002a, 2002b). First, a knowledge engineer works with a subject matter expert to develop an ontology for the application domain. They use the ontology import module (to extract relevant ontology elements from existing knowledge repositories) as well as the various ontology editors and browsers of Disciple-RKF. Figure 5 shows the interfaces of three of Disciple's ontology browsers: at top is the association browser (which displays an object and its relationships with other objects), at bottom left is the tree browser (which displays the hierarchical relationships between the objects in a tree structure), and at bottom right is the hierarchical browser (which displays the hierarchical relationships between the objects in a graph structure).

The result of this knowledge base development phase is an object ontology which is complete enough to be used as a generalization hierarchy for learning, allowing the expert to train the Disciple agent on how to solve problems, with limited assistance from a knowledge engineer. The expert formulates a specific problem solving task and shows the agent the corresponding problem solving steps, helping the agent to understand them. Each problem solving step indicated by the expert consists of a task to be reduced, a question related to that task, the answer to the question, and one or several subtasks or solutions that reduce the task. The top part of Figure 6 shows a fragment of a reasoning tree (generated by the interactive problem solver) consisting of a sequence of task reduction steps. The bottom part shows the interface of the Modeling Assistant that helps the expert express how to reduce the task from the bottom of the reasoning tree. This assistant may suggest the question to be asked, the answer to the question,

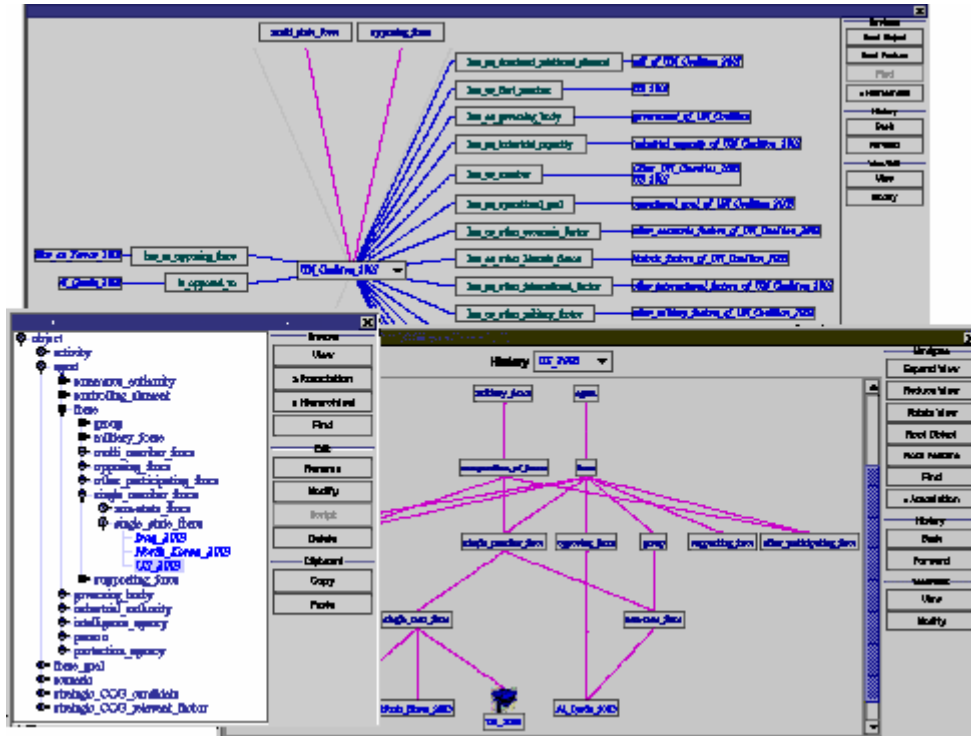


Figure 5: Three ontology browsers of Disciple-RKF.

and the knowledge base elements the expert may be referring to, since the expert's input is in natural language.

Each problem solving step indicated by the expert is an example from which the agent learns a general problem solving rule. First, the expert and Disciple collaborate in formalizing the tasks from the expert's examples and Disciple learns general task patterns. Next the expert helps Disciple find an explanation of why the task reduction step is correct and Disciple learns a general task reduction rule. The top part of Figure 7 shows the interface of the task formalization module, with the tasks in natural language in the middle and their formalizations in the right. The bottom part of Figure 7 shows the explanation generation and selection interface. The task reduction rule learned from the example specified in the modeling assistant (see bottom of Figure 6) is shown in the rule viewer from Figure 8.

As Disciple learns new rules from the expert, the interaction between the expert and Disciple evolves from a teacher-student interaction toward an interaction where both collaborate in solving a problem. During this mixed-initiative problem solving phase, Disciple learns not only from the contributions of the expert, but also from its own successful or unsuccessful problem solving attempts, which leads to the refinement of the learned rules. At the same time, Disciple extends the object ontology with new objects and features.



Figure 6: The reasoning tree and the modeling assistant interfaces.

Copies of Disciple agents may also be trained in parallel by different experts. In this case, the individual knowledge bases have to be merged into an integrated knowledge base, as discussed in section 8.

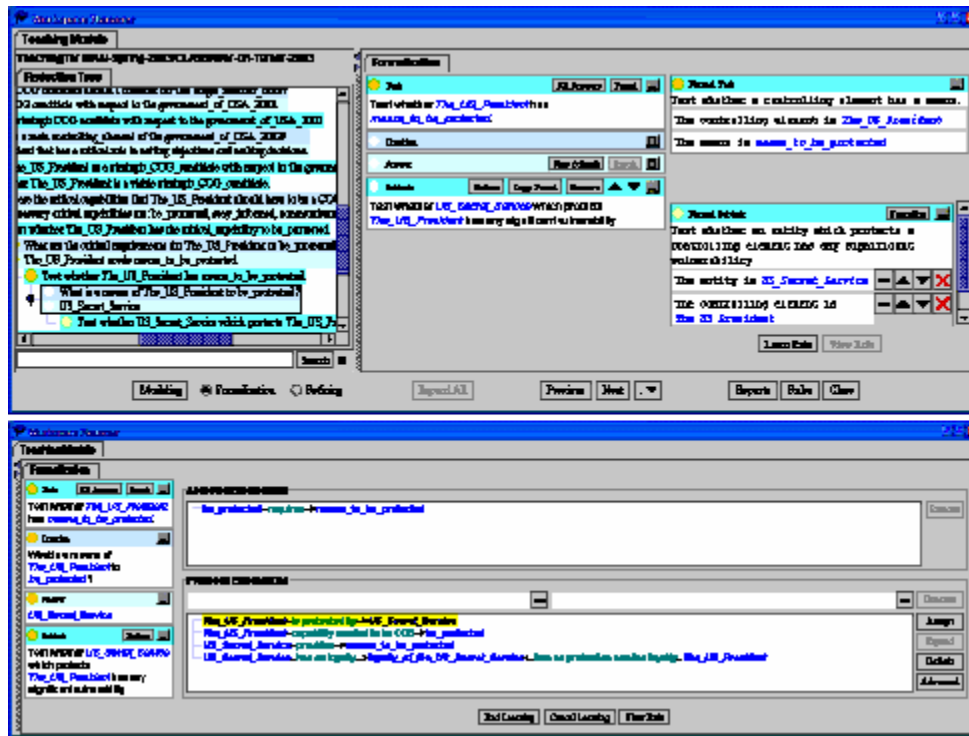


Figure 7: Task formation and example explanation.

6. Use of Disciple-COG in the “Case Studies of Center of Gravity Analysis” Courses: Scenario Elicitation Experiments

Successive versions of the customized Disciple-RKF/COG agent were used in the “Case Studies in Center of Gravity Analysis” course at the Army War College since 2001, becoming part of its regular syllabus. Figure 9 shows how the agent was used. First Disciple was taught how to analyze a scenario, based on the expertise of the course’s instructor. The students then used Disciple as an intelligent assistant that helps them to develop a center of gravity analysis of a war scenario. Each session of this course was an experiment in scenario elicitation from subject matter experts. These experiments demonstrated that the Disciple approach can be used to develop agents that have been found to be useful for a complex military domain (Tecuci et al., 2004c).

7. Use of Disciple-RKF/COG in the “Military Applications of Artificial Intelligence” Courses: Agent Teaching Experiments

Figure 10 shows the use of Disciple in the “589jw Military Applications of Artificial Intelligence” at the US Army War College. In this course, the students teach personal Disciple agents their own expertise in center of gravity analysis and then evaluate both the developed agents and the development process. In the 2001 experiment, the students used historic scenarios with state actors (such as Okinawa 1943) to teach personal Disciple agents how to identify center of gravity candidates. In the 2002 experiment, the students used historic scenarios and a hypothetical scenario with state actors to teach personal Disciple agents how to identify center of

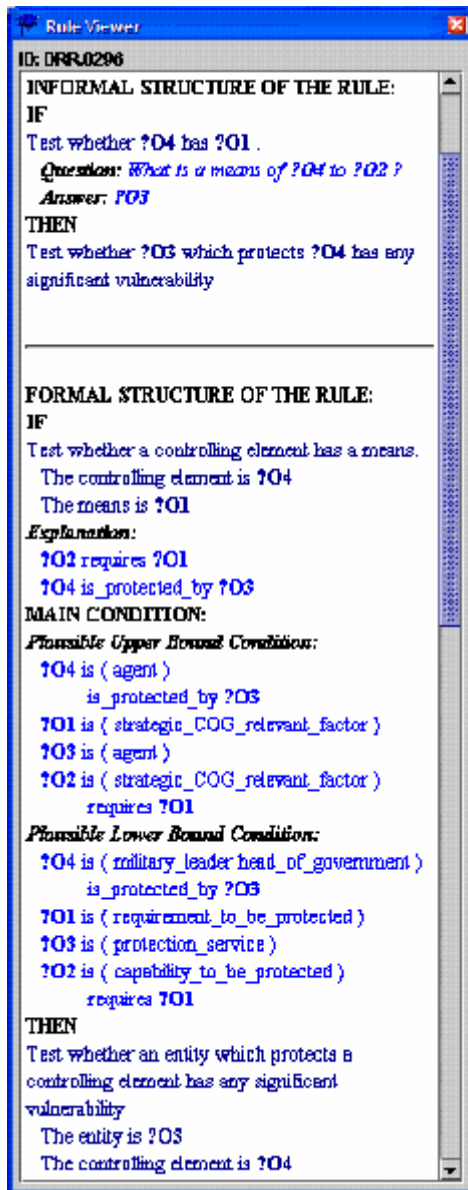


Figure 8: Learned Rule




Figure 9: Training and using the Disciple agent (the COG course).

a domain analysis and ontology development was performed, by involving all the subject matter experts. This considered the example of testing whether Saddam Hussein, in the Iraq 2003 scenario, has all the required critical capabilities to be the center of gravity for Iraq. Based on this domain analysis, the ontology of Disciple-RKF was extended with the definition of 37 new concepts and features identified with the help of the subject matter experts.

gravity candidates and to eliminate those candidates that do not pass certain tests. In the 2003 experiment, the students used historic, current and hypothetical scenarios, with both state and non-state actors, to teach personal Disciple agents how to test center of gravity candidates based on the concepts of critical capabilities, critical requirements and critical vulnerabilities. A total of 38 US and international officers from all the military branches and the national reserve have attended these courses. At the end of these three experiments, 10 of them strongly agreed, 20 agreed, 7 were neutral and only one disagreed with the statement “*I think that a subject matter expert can use Disciple to build an agent, with limited assistance from a knowledge engineer.*” This result shows that significant progress has been made in developing the technology that will allow subject matter experts to build their own intelligent assistants (Tecuci et al., 2004a).

8. Experiment of Parallel Knowledge Base Development by Subject Matter Experts

The Spring 2003 session of the “Military Applications of Artificial Intelligence” course included an experiment in parallel knowledge base development by subject matter experts, which is illustrated in Figure 11 (Tecuci et al., 2004b). Before starting the experiment, the Disciple-RKF agent was trained to identify leaders as center of gravity candidates. The knowledge base of this agent contained the definitions of 432 concepts and features, 29 tasks and 18 task reduction rules. However, the agent had no knowledge of how to test the identified candidates. Then

The 13 subject matter experts were grouped into five teams (of 2 or 3 experts each), and each team was given a copy of the extended Disciple-RKF agent. After that, each team trained its agent to test whether a leader has one or two critical capabilities, as indicated in Figure 1. For instance, Team 1 trained its agent how to test whether a leader has the critical capabilities of staying informed and being irreplaceable. The training was done based on three scenarios (Iraq 2003, Arab-Israeli 1973, and War on Terror 2003), the experts teaching Disciple-RKF how to test each strategic leader from these scenarios. As a result of the training performed by the experts, the knowledge base of each Disciple-RKF agent was extended with new features, tasks, and rules, as indicated in Figure 11. For instance, the knowledge base of the agent trained by Team 1 was extended with 5 features, 10 tasks and 10 rules. The average training time per team was 5 hours and 28 minutes, and the average rule learning rate per team was 3.53 rules/hour. This included the time spent in all the agent training activities (i.e., scenario specification, modeling expert's reasoning, task formalization, rule learning, problem solving, and rule refinement).

After the training of the 5 Disciple-RKF agents, their knowledge bases were merged by a knowledge engineer, who used the knowledge base merging tool of Disciple-RKF. The knowledge engineer also performed a general testing of the integrated knowledge base, in which the acquired 10 new features, 102 new tasks, and 99 new rules were included. During this process two semantically equivalent features were unified, 4 rules were deleted, and 12 other rules were refined by the knowledge engineer. The other 8 features and 83 rules acquired from the experts were not changed. Most of the modifications were done to remove rule redundancies, or to specialize overly general rules.




Figure 10: Teaching and testing a personal Disciple agent (the MAAI courage).

Next, each expert team tested the integrated agent on a new scenario (North Korea 2003), and was asked to judge the correctness of each reasoning step performed by the agent, but only for the capabilities for which that SME team performed the training of the agent. The result was 98.15% correctness.

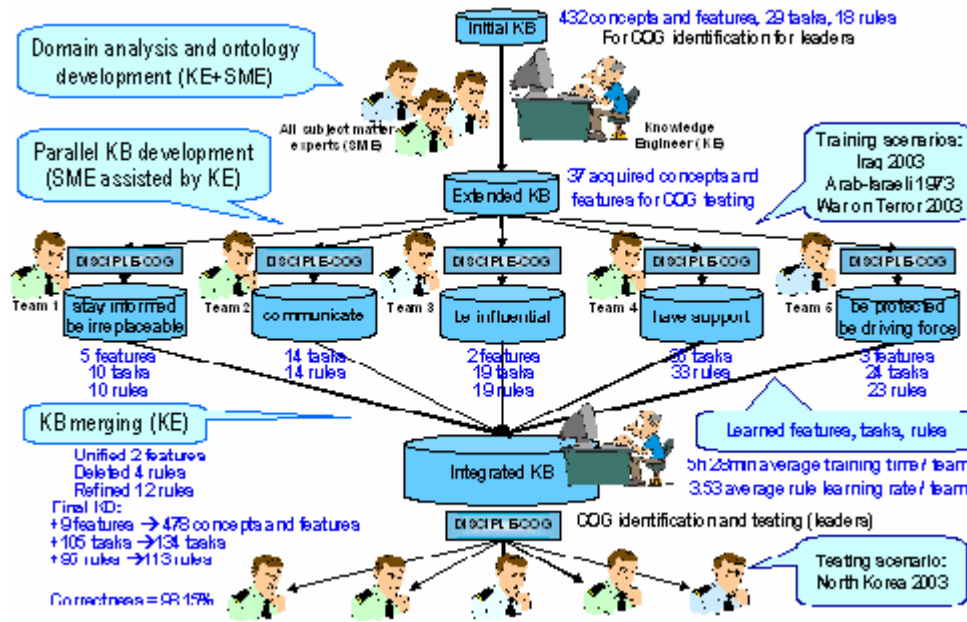


Figure 11: Experiment of rapid knowledge base development by subject matter experts.

9. Conclusions and Future Research Directions

The experiment described in section 8 and illustrated in Figure 11 is the first of this kind ever performed, showing Disciple's capability for rapid and parallel development of knowledge bases by subject matter experts, with limited assistance from knowledge engineers, and for the integration of the developed knowledge bases into a functioning agent. However, while significant progress has been made in developing a capability of rapid knowledge formation by subject matter experts, much work remains to be done to improve the developed methods. For instance, while the subject matter expert has an increased role and independence in the agent development process, the knowledge engineer still has a critical role to play. The knowledge engineer has to assure the development of a fairly complete and correct object ontology. The knowledge engineer also has to develop a generic modeling of the expert's problem solving process based on the task reduction paradigm. Even guided by this generic modeling, and using natural language, the subject matter expert has difficulties in expressing his reasoning process. Therefore more work is needed to develop methods for helping the expert in this task, along the path opened by the Modeling Advisor. The performed experimentations have also revealed that the mixed-initiative reasoning methods of Disciple-RKF could be significantly empowered by developing the natural language processing capabilities of the system.

The development of Disciple-RKF has emphasized initial knowledge acquisition from a subject matter expert, which results in an initial knowledge base. Further work is required to develop knowledge acquisition and learning methods for knowledge base extension and refinement, as well as the integration, validation, and maintenance of the knowledge acquired from different subject matter experts. In particular, the Disciple approach can naturally be extended with methods and tools for:

- Knowledge acquisition for knowledge base refinement (including rule refinement and ontology learning) using mixed-initiative and multistrategy techniques that will exploit the complementarity between human and automated reasoning, and between several learning strategies (such as learning from examples, from explanations, by analogy, by abduction, or by abstraction).
- Acquisition of meta-rules that will capture an expert's rational for choosing among different ways of performing a problem solving task.
- Knowledge bases integration, validation, and maintenance (which will address problems such as ontology merging, inconsistencies within a knowledge piece, redundancies and inconsistencies among knowledge pieces acquired from different experts, refinement and reorganization of the object ontology for increased performance, restructuring and refinement of the acquired problem solving rules, etc.).

10. Appendix

10.1 Publications describing the results of this project

Barbulescu M., Balan G., Boicu M., Tecuci G., "Rapid Development of Large Knowledge Bases," in *Proceedings of the 2003 IEEE International Conference on Systems, Man & Cybernetics*, October 5–8, Washington D.C., 2003.

Boicu C., Tecuci G., "Mixed-Initiative Ontology Learning," in *Proceedings of the 2004 International conference on Artificial Intelligence*, IC-AI'2004, June 21 - 24, 2004, Monte Carlo Resort, Las Vegas, Nevada, USA. IEEE Computer Society, Los Alamitos, California, 2004.

Boicu C., Tecuci G., Boicu M., "Mixed-Initiative Exception-Based Learning for Knowledge Base Refinement," in *Proceedings of the Eighteenth National Conference of Artificial Intelligence and the Fourteenth Conference on Innovative Applications of Artificial Intelligence*, AAAI-02/IAAI-02, pp. 947-948, Edmonton, Alberta, Canada, AAAI Press/The MIT Press, 2002.

Boicu C., Tecuci G., Boicu M., Marcu D., "Improving the Representation Space through Exception-Based Learning," in *Proceedings of the 16th International FLAIRS Conference (FLAIRS-2003), Special Track on Machine Learning*, May 2003, Key West, Florida. AAAI Press, Menlo Park, CA, pp. 336-340, 2003.

Boicu M., "Modeling and Learning with Incomplete Knowledge," *PhD Dissertation in Information Technology*, Learning Agents Laboratory, School of Information Technology and Engineering, George Mason University, 2002.

Boicu M., Marcu D., Bowman M., and Tecuci G., "A Mixed-Initiative Approach to Teaching Agents to Do Things," in *Proceedings of the Symposium on Learning How to Do Things*, The 2000 AAAI Fall Symposium Series, North Falmouth, Massachusetts, November 3-5, 2000.

Boicu M., Tecuci G. and Stanescu B., "Mixed-Initiative Agent Teaching and Learning," in *Proceedings of the 2001 International Conference on Artificial Intelligence*, IC-AI'2001, pp.

122-128. June 25-28, 2001, Monte Carlo Resort, Las Vegas, Nevada, USA, IEEE Computer Society, Los Alamitos, California, 2001a.

Boicu M., Tecuci G., Stanescu B., Balan G.C. and Popovici E., "Ontologies and the Knowledge Acquisition Bottleneck," in *Proceedings of IJCAI-2001 Workshop on Ontologies and Information Sharing*, pp. 9-18. Seattle, Washington, August 2001, AAAI Press, Menlo Park, California, 2001b.

Boicu M., Tecuci G., Stanescu B., Marcu M. and Cascaval C.E., "Automatic Knowledge Acquisition from Subject Matter Experts," in *Proceedings of the Thirteenth International Conference on Tools with Artificial Intelligence (ICTAI)*, pp. 69-78. 7-9 November 2001, Dallas, Texas, IEEE Computer Society, Los Alamitos, California, 2001c.

Boicu M., Tecuci G., Marcu D., Stanescu B., Boicu C., Balan C., Barbulescu M. and Hao X., "Disciple-RKF/COG: Agent Teaching by Subject Matter Experts," in *Proceedings of the Eighteenth National Conference of Artificial Intelligence and the Fourteenth Conference on Innovative Applications of Artificial Intelligence*, AAAI-02/IAAI-02, pp. 992-993, Edmonton, Alberta, Canada, AAAI Press/The MIT Press, 2002.

Boicu C., Tecuci G., Boicu M., Marcu D., "Improving the Representation Space through Exception-Based Learning," in *Proceedings of the 16th International FLAIRS Conference (FLAIRS-2003)*, Special Track on Machine Learning, May 2003, Key West, Florida. AAAI Press, Menlo Park, CA, 2003a.

Boicu M., Tecuci G., Marcu D., Boicu C., Stanescu B., "Mixed-Initiative Control for Teaching and Learning in Disciple," in *Proceedings of IJCAI-2003 Workshop on Mixed-Initiative Intelligent Systems*, Acapulco, Mexico, August, AAAI Press, Menlo Park, CA, 2003b.

Boicu M., Tecuci G., Stanescu B., Marcu D., Barbulescu M., Boicu C., "Design Principles for Learning Agents," in *Proceedings of AAAI-2004 Workshop on Intelligent Agent Architectures: Combining the Strengths of Software Engineering and Cognitive Systems*, July 26, San Jose, AAAI Press, Menlo Park, CA, 2004.

Bowman M., "A Methodology for Modeling Expert Knowledge that Supports Teaching-Based Development of Agents," *PhD Dissertation in Information Technology*, School of Information Technology and Engineering, George Mason University, Fairfax Virginia, 2002.

Bowman M., Lopez A.M. Jr., Donlon J., and Tecuci G., "Teaching Intelligent Agents: Software Design Methodology," *CROSSTALK: The Journal of Defense Software Engineering*, Vol. 14, No.6, pp. 10-14, 2001a.

Bowman M., Lopez, A.M. Jr., and Tecuci G., "Ontology Development for Military Applications," *Proceedings of the South Eastern Regional ACM Conference*, Atlanta, GA, May 2001b.

Bowman M., Tecuci G., and Boicu M., "Intelligent Agents in the Command Post," in *Military Review*, March-April, pp. 46-53, US Army CGSG, Fort Leavenworth, Kansas, 2001c.

Bowman M., Tecuci G., and Boicu M., "Intelligent Agents: Tools for the Command Post and Commander," in *Army Acquisition, Logistics and Technology (ALT)*, pp.34-35, Nov-Dec 2001d.

Bowman M., Tecuci G., and Ceruti M., "Application of Disciple to Decision Making in Complex and Constrained Environments," in *Proceedings of the 2001 IEEE Systems, Man, and Cybernetics Conference*, pp. 2932-2940, October 2001e.

Cleckner W., Tecuci G., Comello J., "Instructable Agents for Strategic Center of Gravity Analysis," in *Proceedings of the Army Operations Research Symposium*, AORS-2001, 5pg, October 2001.

Cleckner W., Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Comello J., Lopez A., Donlon J., "Teaching Strategic Center of Gravity Analysis Through Learning Agent Technologies," in *Proceedings of the Military Operations Research Symposium*, 6pg, Ft. Leavenworth, June 2002.

Fowler C.W., Cleckner W., Lopez A., Tecuci G., Comello J., "Strategic Center of Gravity Analysis for Military Operations other than War through Learning Agent Technologies," in *Proceedings of the Army Operations Research Symposium*, AORS-2002, October 2002.

Lopez, Jr. A.M., Comello J.J., Bowman M., Tecuci G., Donlon J.J., "Center of Gravity Meets Learning Agent Technology," in *Military Review*, November-December, *US Army CGSG*, Fort Leavenworth, Kansas, 2002.

Marcu D., Boicu M., and Tecuci G., "User-Agent Interactions in Mixed-Initiative Learning," in *Proceedings of the 14th International FLAIRS Conference (FLAIRS-2001)*, Special Track on Machine Learning, May 2001, Key West, Florida. AAAI Press, Menlo Park, CA, 2001.

Shyr P., Tecuci G. and Boicu M., "Evaluation of Mixed-Initiative Knowledge Base Development Methods and Tools," in *Proceedings of IJCAI-2001 Workshop on Empirical Methods in AI*, pp. 47-53, Seattle, Washington, August. AAAI Press, Menlo Park, CA, 2001.

Stanescu B., Boicu C., Balan G., Barbulescu M., Boicu M., Tecuci G., "Ontologies for Learning Agents: Problems, Solutions and Directions," in *Proceedings of the IJCAI-03 Workshop on Workshop on Ontologies and Distributed Systems*, Acapulco, Mexico, August, AAAI Press, Menlo Park, CA, 2003.

Tecuci G., Boicu M., "Military Applications of the Disciple Learning Agent," Chapter 8 in Jain L, (ed.), *Advanced Information Systems in Defense and Related Applications*, pp. 337-376, Springer Verlag, 2002.

Tecuci G., Boicu M., Bowman M., and Marcu D., with a commentary by Burke M., "An Innovative Application from the DARPA Knowledge Bases Programs: Rapid Development of a

High Performance Knowledge Base for Course of Action Critiquing,” in *AI Magazine*, 22, 2, 2001, pp. 43-61. AAAI Press, Menlo Park, California, 2001.

Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Comello J., “Training and Using Disciple Agents: A Case Study in the Military Center of Gravity Analysis Domain,” in *AI Magazine*, 24, 4, 2002, pp.51-68, AAAI Press, Menlo Park, California, 2002a.

Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Comello J., Lopez A., Donlon J., Cleckner W., “Development and Deployment of a Disciple Agent for Center of Gravity Analysis,” in *Proceedings of the Eighteenth National Conference of Artificial Intelligence and the Fourteenth Conference on Innovative Applications of Artificial Intelligence*, AAAI-02/IAAI-02, pp. 853-860, Edmonton, Alberta, Canada, AAAI Press/The MIT Press, 2002b. (Deployed Application Award).

Tecuci G., Aha D., Boicu M., Cox M., Ferguson G., and Tate A. (eds), *Proceedings of the IJCAI-03 Workshop on Mixed-Initiative Intelligent Systems*, Acapulco, Mexico, August 2003a.

Tecuci G., Boicu M., Marcu D., “Toward a Disciple-based Mixed-Initiative Cognitive Assistant,” in *Proceedings of IJCAI-2003 Workshop on Mixed-Initiative Intelligent Systems*, Acapulco, Mexico, August, AAAI Press, Menlo Park, CA, 2003b.

Tecuci G., Boicu M., Boicu C., Marcu D., Stanescu B., Barbulescu M., The Disciple-RKF Learning and Reasoning Agent, *Research Report submitted for publication*, Learning Agents Center, George Mason University, September 2004a.

Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Barbulescu M., "Parallel Knowledge Base Development by Subject Matter Experts" in *Proceedings of the 14th International Conference on Knowledge Engineering and Knowledge Management*, EKAW 2004, 5-8th October 2004 - Whittlebury Hall, Northamptonshire, UK, Springer-Verlag, 2004b.

Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Barbulescu M., "A University Research Group Experience with Deploying an Artificial Intelligence Application in the Center of Gravity Analysis Domain," in *Proceedings of AAAI-2004 Workshop on Fielding Applications of Artificial Intelligence*, July 25, San Jose, AAAI Press, Menlo Park, CA, 2004c.

Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Barbulescu M., "A Learning Agent Shell for Building Knowledge-Based Agents," in *Proceedings of the Technology Demonstration Session of the 14th International Conference on Knowledge Engineering and Knowledge Management*, EKAW 2004, 5-8th October 2004 - Whittlebury Hall, Northamptonshire, UK, 2004d.